arXiv:1410.0827v1 [stat.ME] 3 Oct 2014

Multiscale Bernstein polynomials for densities

Antonio Canale* and David B. Dunson†

October 6, 2014

Abstract

Our focus is on constructing a multiscale nonparametric prior for densities. The Bayes density estimation literature is dominated by single scale methods, with the exception of Polya trees, which favor overly-spiky densities even when the truth is smooth. We propose a multiscale Bernstein polynomial family of priors, which produce smooth realizations that do not rely on hard partitioning of the support. At each level in an infinitely-deep binary tree, we place a beta dictionary density; within a scale the densities are equivalent to Bernstein polynomials. Using a stick-breaking characterization, stochastically decreasing weights are allocated to the finer scale dictionary elements. A slice sampler is used for posterior computation, and properties are described. The method characterizes densities with locally-varying smoothness, and can produce a sequence of coarse to fine density estimates. An extension for Bayesian testing of group differences is introduced and applied to DNA methylation array data.

Keywords: Density estimation; Multiresolution; Multiscale clustering; Multiscale testing; Nonparametric Bayes; Polya tree; Stick-breaking; Wavelets

*Dipartimento di Scienze Economico-Sociali e Matematico-Statistiche, Università degli studi di Torino e Collegio Carlo Alberto, Torino, Italy antonio.canale@unito.it

†Department of Statistical Science, Duke University, Durham, NC dunson@duke.edu

1 Introduction

Multiscale estimators have well known advantages, including the ability to characterize abrupt local changes and to provide a compressed estimate to a desired level of resolution. Such advantages have lead to enormous popularity of wavelets, which are routinely used in signal and image processing, and have had attention in the literature on density estimation. Donoho et al. (1996) developed a wavelet thresholding approach for density estimation, which has minimax optimality properties, and there is a literature developing modifications for deconvolution problems (Pensky and Vidakovic 1999), censored data (Niu 2012), time series (Garcia-Trevino and Barria 2012) and other settings. Locke and Peter (2013) proposed an approach, which can better characterize local symmetry and other features commonly observed in practice, using multiwavelets. Chen et al. (2012) instead use geometric multiresolution analysis methods related to wavelets to obtain estimates of high-dimensional distributions having low-dimensional support.

Although there is a rich Bayesian literature on multiscale function estimation (Abramovich et al. 1998; Clyde et al. 1998; Clyde and George 2000; Wang et al. 2007), there has been limited consideration of Bayesian multiscale density estimation. Popular methods for Bayes density estimation rely on kernel mixtures. For example, Dirichlet process mixtures are applied routinely. By using location-scale mixtures, one can accommodate varying smoothness, with the density being flat in certain regions and concentrated in others. However, Dirichlet processes lack the appealing multiscale structure. Polya trees provide a multiscale alternative (Mauldin et al. 1992; Lavine 1992a,b), but have practical disadvantages. They tend to produce highly spiky density estimates even when the true density is smooth, and have sensitivity to a pre-specified partition sequence. This sensitivity can be ameliorated by mixing Polya trees (Hanson and Johnson 2002), but at the expense of more difficult computation.

Our focus is on developing a new approach for Bayesian multiscale density estimation, which inherits many of the advantages of Dirichlet process mixtures while avoiding the key disadvantages of Polya trees. We want a framework that is easily computable, has desirable multiscale approximation properties, allows centering on an initial guess at the density, and can be extended in a straightforward manner to include covariates and allow embedding within larger models. We accomplish this using a multiscale extension of mixtures of Bernstein polynomials (Petrone 1999a,b), which have been shown to have appealing asymptotic properties in the single scale case (Petrone and Wasserman 2002; Ghosal 2001).

In the next section, our multiscale prior for densities is introduced and properties are discussed. Section 3 introduces posterior computation via a slice sampling algorithm. In Section 4 the performance of the method in terms of density estimation is evaluated via a simulation study. Section 5 discusses generalizations, with particular emphasis on Bayesian multiscale inferences on differences between groups. Section 6 applies the method to a DNA methylation array dataset on breast cancer, and Section 7 concludes. Proofs and computational details are reported in the Appendix.

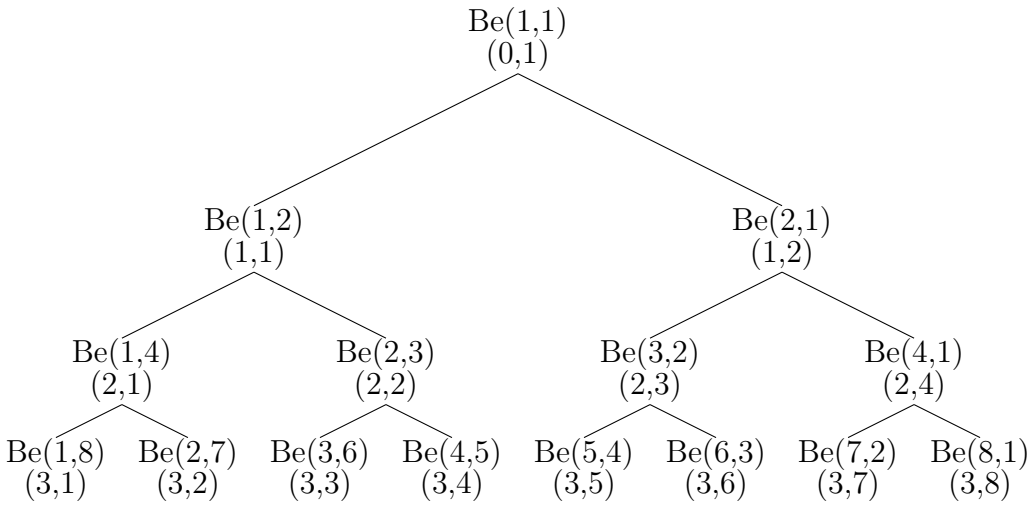

Figure 1: Binary tree with beta kernels at each node (s, h), where s is the scale level and h is the index within the scale

2 Multiscale priors for densities

2.1 Proposed model

Let $x \in \mathcal{X} \subset \mathbb{R}$ be a random variable having density g with respect to Lebesgue measure. Assume that g_0 is a prior guess for g, with G_0 and G_0^{-1} the corresponding cumulative distribution function (CDF) and inverse CDF, respectively. We induce a prior $g \sim \Pi$ centered on g_0 through a prior for the density f of $y = G_0(x) \in (0, 1)$. The CDFs F and G corresponding to the densities f and g, respectively, have the following relationship

$$G(x) = F\{G_0(x)\}, x \in \mathcal{X}, \quad F(y) = G\{G_0^{-1}(y)\}, y \in (0, 1). \tag{1}$$

We assume that f follows a multiscale mixture of Bernstein polynomials,

$$f(y) = \sum_{s=0}^{\infty} \sum_{h=1}^{2^s} \pi_{s,h} \text{Be}(y; h, 2^s - h + 1), \tag{2}$$

where Be(a, b) denotes the beta density with mean $a/(a + b)$, and $\{\pi_{s,h}\}$ are random weights drawn from a suitable stochastic process. We introduce an infinite sequence of scales $s = 0, 1, \ldots, \infty$. At scale s, we include 2^s Bernstein polynomial basis densities. The framework can be represented as a binary tree in which each layer is indexed by a scale and each node is a suitable beta density. For example, at the root node, we have the Be(1,1) density which generates two daughters Be(1,2) and Be(2,1) and so on. In general, let s denote the scale and h the polynomial within the scale. The node (s, h) in the tree is related to the Be$(h, 2^s - h + 1)$ density. A cartoon of the binary tree is reported in Figure **??**.

A prior measure for the multiscale mixture (2) is obtained by specifying a stochastic process for the infinite dimensional set of weights $\{\pi_{s,h}\}$. To this end we introduce, for

each scale s and node h within the scale, independent random variables

$$S_{s,h} \sim \text{Be}(1, a), \quad R_{s,h} \sim \text{Be}(b, b), \tag{3}$$

corresponding to the probability of stopping and taking the right path conditionally on not stopping, respectively. Define the weights as

$$\pi_{s,h} = S_{s,h} \prod_{r<s} (1 - S_{r,g_{shr}}) T_{shr} \tag{4}$$

where $g_{shr} = \lceil h/2^{s-r} \rceil$ is the node traveled through at scale r on the way to node h at scale s, $T_{shr} = R_{r,g_{shr}}$ if $(r+1, g_{shr+1})$ is the right daughter of node (r, g_{shr}), and $T_{shr} = 1 - R_{r,g_{shr}}$ if $(r+1, g_{shr+1})$ is the left daughter of (r, g_{shr}). For binary trees, there is a unique path leading from the root node to node (s, h), and $\mathcal{T}$ denotes the infinite deep binary tree of the weights (4). We refer to the prior resulting from (2)–(4) as a multiscale Bernstein polynomial (msBP) prior and we write $f \sim \text{msBP}(a, b)$. The choice for the hyperparameters are discussed in the next section.

The infinite tree of probability weights is generated from a generalization of the stick-breaking process representation of the Dirichlet process (Sethuraman 1994). Each time the stick is broken, it is consequently randomly divided in two parts (one for the probability of going right, the remainder for the probability of going left) before the next break. An alternative treed stick-breaking process is proposed by Adams et al. (2010) where a first stick-breaking process defines the vertical growth of an infinitely wide tree and a second puts weights on the infinite number of descendant nodes.

Sampling a random variable y from a random density, which is generated from a msBP prior, can be described as follows. At node (s, h), generate a random probability $S_{s,h} \sim \text{Be}(1, a)$ corresponding to the probability of stopping at that node given you passed through that node, and $R_{s,h} \sim \text{Be}(b, b)$ corresponding to the probability of taking the right path in the tree in moving to the next finer scale given you did not stop at node (s, h). Conditionally on being at the node (s, h) we assume that $y \sim \text{Be}(y; h, 2^s - h + 1)$. Algorithm 1 describes how to generate y from an msBP density.

Algorithm 1 Generating a draw from a random density having an msBP prior

```
loop = TRUE;
s = 0, h = 1;
while loop do
   let loop = FALSE with probability S_{s,h}.
   if loop then
      with probability R_{s,h}, let h = 2h
      with probability 1 − R_{s,h}, let h = 2h − 1
   end if
end while
generate y ∼ Be(h, 2^s − h + 1).
```

2.2 Basic properties

In this section we study basic properties of the proposed prior. A first requirement is that the construction leads to a meaningful sequence of weights. The next lemma

shows that the random weights on each node of the infinitely deep tree sum to one almost surely.

Lemma 1. *Let $\pi_{s,h}$ be an infinite sequence of weights defined as in* (3)–(4)*. Then,*

$$\sum_{s=0}^{\infty}\sum_{h=1}^{2^s}\pi_{s,h} = 1 \tag{5}$$

almost surely for any $a, b > 0$.

The total weight placed on a scale s is controlled by the prior for $S_{s,h}$. The expected probability allocated to node h at scale s can be expressed as

$$\begin{aligned}\mathrm{E}(\pi_{s,h}) &= \mathrm{E}\left\{S_s\prod_{l=0}^{s-1}(1-S_l)\prod_{l=1}^{s}T_l\right\} \\ &= \left(\frac{1}{1+a}\right)\left(\frac{a}{1+a}\right)^s\left(\frac{1}{2}\right)^s = \frac{1}{1+a}\left(\frac{a}{2+2a}\right)^s,\end{aligned} \tag{6}$$

where we discard the h subscript on $S_l \sim \mathrm{Be}(1, a)$ and $T_l \sim \mathrm{Be}(b, b)$ for ease in notation. This does not impact the calculation because any path taken up to scale s has the same probability *a priori* and the random variables in (3) have the same distribution regardless of the path that is taken. Similarly

$$\mathrm{E}(\pi_{s,h}^2) = \mathrm{E}\left\{S_s^2\prod_{l=0}^{s-1}(1-S_l)^2\prod_{l=1}^{s}T_l^2\right\}, = \frac{2}{(1+a)(2+a)}\left(\frac{a}{2+a}\right)^s\left\{\frac{b+1}{2(2b+1)}\right\}^s.$$

Hence at scale $s = 0$ the variance is $\mathrm{Var}(\pi_{0,1}) = a/\{(2+a)(1+a)^2\}$, while for $s > 0$

$$\mathrm{Var}(\pi_{s,h}) = \frac{2}{(1+a)(2+a)}\left(\frac{a}{2+a}\right)^s\left\{\frac{b+1}{2(2b+1)}\right\}^s - \left\{\frac{1}{1+a}\left(\frac{a}{2+2a}\right)^s\right\}^2. \tag{7}$$

We can additionally verify that our prior for the CDF G is centered on the chosen G_0. Letting $F(A) = \int_A f$, we obtain $E\{F(A)\} = \lambda(A)$, where $\lambda(A)$ is the Lebesgue measure over the set A. Details are reported in the Appendix. Hence, the prior for the density of y is automatically centered on a uniform density on $[0, 1]$. This is the desired behavior as $y \sim \mathrm{Unif}(0, 1)$ with $x = G_0^{-1}(y)$ implies that $x \sim g_0$, which is our prior guess for the observed data density. In addition, from (1), $\mathrm{E}\{F(y)\} = y$ implies

$$\mathrm{E}[G\{G_0^{-1}(y)\}] = y = \mathrm{E}\{G(x)\} = G_0(x),$$

so that the prior expectation for the CDF G is G_0 as desired.

From equation (6) and (7), the hyperparameter a controls the decline in probabilities over scales. In general, letting $S^{(i)}$ denote the scale at which the ith observation falls, we have

$$E(S^{(i)}) = \sum_{s=0}^{\infty} s\frac{1}{1+a}\left(\frac{a}{2+2a}\right)^s = a.$$

Hence, the value of a is the expected scale from which observations are drawn. For small a, high probability is placed on coarse scales, leading to smoother densities, with

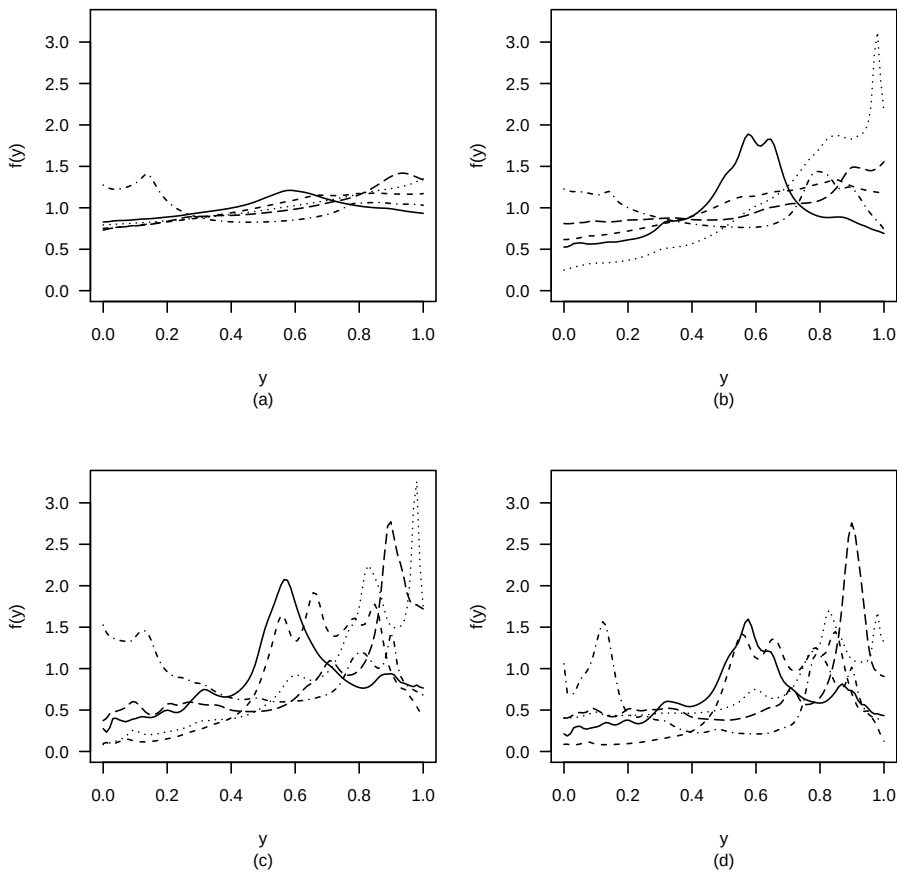

Figure 2: Five realizations from an msBP prior with $b = 1$ and (a) $a = 1$, (b) $a = 2$, (c) $a = 5$, and (d) $a = 10$.

$a \to 0$ inducing $\pi_{0,1} = 1$ and hence $f(y)$ uniform. As a increases, finer scale densities will be weighted higher, leading to spiker realizations. To illustrate this, Figure 2 shows realizations from the prior for different a values. To better isolate the contribution of the a hyperparameter, we fixed the realizations of $R_{s,h} \sim \text{Be}(1, 1)$ for all subplots.

An appealing aspect of the proposed formulation is that individuals sampled from a distribution that is assigned an msBP prior are allocated to clusters in a multiscale fashion. In particular, two individuals having similar observations may have the same cluster allocation up to some scale s, but perhaps are not clustered on finer scales. Clustering is intrinsically a scale dependent notion, and our model is the first to our knowledge to formalize multiscale clustering in a model based probabilistic manner. Under the above structure, the probability that two individuals i and i' are assigned

to the same scale s cluster is one for $s = 0$ and for $s > 0$, is equal to

$$2^s \mathrm{E}\Big\{ \prod_{l=0}^{s-1} (1 - S_l)^2 T_l^2 \Big\}$$
$$= 2^s \prod_{l=0}^{s-1} \mathrm{E}(\overline{S}_l^2)\mathrm{E}(T_l^2) = 2^s \left(\frac{a}{a+2} \right)^s \left(\frac{1}{2} \right)^s \left(\frac{b+1}{2b+1} \right)^s = \left\{ \left(\frac{a}{a+2} \right) \left(\frac{b+1}{2b+1} \right) \right\}^s .$$

This is derived by calculating the expected probability that two individuals travel though node h at scale s and multiplying by the number of nodes in scale s. This form is intuitive. As $b \to 0$, the $\mathrm{Be}(b, b)$ density degenerates to $0.5\delta_0 + 0.5\delta_1$, so that variability among subjects in the chosen paths through the tree decreases and all subjects take a common path chosen completely at random via unbiased coin flips at each node. In such a limiting case, $(b+1)/(2b+1) \to 1$ and the probability of clustering subjects at scale s is simply the probability of surviving to that scale and not being allocated to a coarser scale component. At the other extreme, as $b \to \infty$ each subject independently flips an unbiased coin in deciding to go right or left at each node of the tree, and $(b + 1)/(2b + 1) \to 1/2$. Hyperpriors can be chosen for a and b to allow the data to inform about these tuning parameters; we find that choosing a hyperprior for a is particularly important, with $b = 1$ as a default.

Approximations of the msBP process can be obtained fixing an upper bound s for the depth of the tree. The truncation is applied by pruning $\mathcal{T}$ at scale s, setting $S_{s,h} = 1$ for each $h = 1, \ldots, 2^s$ as done in Ishwaran and James (2001) and related works in the single scale case. We denote the scale s approximation as

$$f^s(y) = \sum_{l=0}^{s} \sum_{h=1}^{2^l} \tilde{\pi}_{l,h} \mathrm{Be}(y; h, 2^l - h + 1), \tag{8}$$

with $\tilde{\pi}_{l,h}$ identical to $\pi_{l,h}$ except that we set all the stopping probabilities at scale s equal to one to ensure that the weights sum to one and that $f^s(y)$ is a valid probability density on $\mathcal{Y} = [0, 1]$. Let $\mathcal{T}^s$ denote the pruned binary tree of weights. It is interesting to study the accuracy of the approximation of $f^s(y)$ to $f(y)$ as the scale s changes under different metrics. For example, using the total variation distance,

$$d_{TV}(P_s, P) = \sup_{B \in \mathcal{B}} |P^s(B) - P(B)|$$
$$= \sup_{B \in \mathcal{B}} \left| \sum_{h=1}^{2^s} \tilde{\pi}_{s,h} \mathrm{Be}(B; h, 2^s - h + 1) - \sum_{l=s}^{\infty} \sum_{h=1}^{2^l} \pi_{l,h} \mathrm{Be}(B; h, 2^l - h + 1) \right|, \tag{9}$$

where $P^s(B) = \int_B f^s(y)dy$ and $P(B) = \int_B f(y)dy$, for all $B \in \mathcal{B}$, denote the probability measures corresponding to densities $f^s(y)$ and $f(y)$, respectively, with $\mathcal{B}$ the Borel σ-algebra of subsets of $\mathcal{Y} = [0, 1]$. The next lemma shows that *a priori* the expected deviation of the truncation approximation P^s from P is zero and the variance is decreasing exponentially with s.

Lemma 2. *The expectation of the total variation distance between $P^s(B)$ and $P(B)$*

is zero and its variance is

$$Var\{d_{TV}(P_s, P)\} = 2\left(\frac{a}{a+1}\right)^s.$$

3 Posterior computation

In this section we demonstrate that a straightforward Markov chain Monte Carlo (MCMC) algorithm can be constructed to perform posterior inference under the msBP prior. The algorithm consists of two primary steps: (i) allocate each observation to a multiscale cluster, conditionally on the current values of the probabilities $\{\pi_{s,h}\}$; (ii) conditionally on the cluster allocations, update the probabilities.

Suppose subject i is assigned to node (s_i, h_i), with s_i the scale and h_i the node within scale. Conditionally on $\{\pi_{s,h}\}$, the posterior probability of subject i belonging to node (s, h) is simply

$$\text{pr}(s_i = s, h_i = h|y_i, \pi_{s,h}) \propto \pi_{s,h}\text{Be}(y; h, 2^s - h + 1).$$

Consider the total mass assigned at scale s, defined as $\pi_s = \sum_{h=1}^{2^s} \pi_{s,h}$, and let $\bar{\pi}_{s,h} = \pi_{s,h}/\pi_s$. Under this notation, we can rewrite (2) as

$$f(y) = \sum_{s=0}^{\infty} \pi_s \sum_{h=1}^{2^s} \bar{\pi}_{s,h}\text{Be}(y; h, 2^s - h + 1).$$

To allocate each subject to a multiscale cluster, we rely on a multiscale modification of the slice sampler of Kalli et al. (2011). Consider the joint density

$$f(y_i, u_i, s_i) \propto \mathbb{1}(u_i < \pi_{s_i}) \sum_{h=1}^{2^{s_i}} \bar{\pi}_{s_i,h}\text{Be}(y_i; h, 2^{s_i} - h + 1).$$

The full conditional posterior distributions are

$$u_i|y_i, s_i \sim U(0, \pi_{s_i}), \tag{10}$$

$$\text{pr}(s_i = s|u_i, y_i) \propto \mathbb{1}(s : \pi_s > u_i) \sum_{h=1}^{2^s} \bar{\pi}_{s,h}\text{Be}(y_i; h, 2^s - h + 1), \tag{11}$$

$$\text{pr}(h_i = h|u_i, y_i, s_i) \propto \bar{\pi}_{s_i,h}\text{Be}(y_i; h, 2^{s_i} - h + 1). \tag{12}$$

Even with an infinite resolution level, equation (11) implies that observations are assigned to a finite number of scales and there are a finite number of probabilities to evaluate. Conditionally on the scale, equation (12) induces a simple multinomial sampling, which allocates a subject to a particular node within that scale. Algorithm 2 summarizes the posterior cluster allocation step. An alternative version of this slice sampler considers the joint density

$$f(y_i, u_i, s_i, h_i) \propto \mathbb{1}(u_i < \pi_{s_i,h_i})\text{Be}(y_i; h_i, 2^{s_i} - h_i + 1),$$

leading to conditional posteriors

$$u_i|y_i, s_i, h_i \sim U(0, \pi_{s_i,h_i}), \quad \text{pr}(s_i = s, h_i = h|u_i, y_i) \propto \mathbb{1}(\pi_{s,h} > u_i)\text{Be}(y_i; h, 2^s - h + 1).$$

In the second version a greater number of probabilities need to be evaluated for each subject. Our experience suggests that the sampler obtained using (10)–(12), summarized in Algorithm 2, is more efficient and converges faster.

Algorithm 2 Multiscale cluster posterior allocation for ith subject

for each scale s **do**
  calculate $\pi_s = \sum_{h=1}^{2^s} \pi_{s,h}$:
end for
simulate $u_i|y_i, s_i \sim U(0, \pi_{s_i})$;
for each scale s **do**
  if $\pi_s > u_i$ **then**
    for $h = 1, \dots 2^s$ **do**
      compute $\bar{\pi}_{s,h} = \pi_{s,h}/\pi_s$
    end for
    compute $\text{pr}(s_i = s|u_i, y_i) \propto \sum_{h=1}^{2^s} \bar{\pi}_{s,h}\text{Be}(y_i; h, 2^s - h + 1)$
  else
    $\text{pr}(s_i = s|u_i, y_i) = 0$;
  end if
end for
sample s_i with probability $\text{pr}(s_i = s|u_i, y_i)$;
sample h_i with probability $\text{pr}(h_i = h|y_i, s_i) \propto \bar{\pi}_{s_i,h}\text{Be}(y_i; h, 2^{s_i} - h + 1)$;

Conditionally on cluster allocations, we sample all the stopping and descending-right probabilities from their full conditional posterior distributions:

$$S_{s,h} \sim \text{Be}(1 + n_{s,h}, a + v_{s,h} - n_{s,h}), \quad R_{s,h} \sim \text{Be}(b + r_{s,h}, b + v_{s,h} - n_{s,h} - r_{s,h}), \quad (13)$$

where $v_{s,h}$ is the number of subjects passing through node (s, h), $n_{s,h}$ is the number of subjects stopping at node (s, h), and $r_{s,h}$ is the number of subjects that continue to the right after passing through node (s, h). Calculation of $v_{s,h}$ and $r_{s,h}$ can be performed via parallel computing due to the binary tree structure, improving efficiency.

If hyperpriors for a and b are assumed, additional sampling steps are required. Assuming $a \sim \text{Ga}(\beta, \gamma)$, its full conditional posterior is

$$a|- \sim \text{Ga}\left(\beta + 2^{s'+1} - 1, \gamma - \sum_{s=0}^{s'}\sum_{h=1}^{2^s} \log(1 - S_{s,h})\right), \quad (14)$$

while if $b \sim \text{Ga}(\delta, \lambda)$ its full conditional posterior is proportional to

$$b^\delta \prod_{s=0}^{s'}\prod_{h=1}^{2^s} \frac{1}{B(b, b)} \exp\left\{-b\left(\lambda \sum_{s=0}^{s'}\sum_{h=1}^{2^s} \log\{R_{s,h}(1 - R_{s,h})\}\right)\right\}, \quad (15)$$

where s' is the maximum occupied scale and $B(p, q)$ is the Beta function. To sample from the latter distribution, a Metropolis-Hastings step is required. The Gibbs sampler iterates the steps outlined in Algorithm 3.

Algorithm 3 Gibbs sampler steps for posterior computation under msBP prior

for $i = 1, \ldots, n$ **do**
  assign observation i to a cluster (s_i, h_i) as in Algorithm 2.
end for
compute $n_{s,h}$ the number of subjects in cluster (h, s) for all occupied clusters;
compute $v_{s,h}$ the number of subjects that pass through node (h, s);
compute $r_{s,h}$ the number of subjects that proceed down to the right at node (h, s);
let s_{MAX} be the maximum occupied scale;
for $s = 0, \ldots, s_{\text{MAX}}$ **do**
  for $h = 1, \ldots, 2^s$ **do**
    update $S_{s,h} \sim \text{Be}(1 + n_{s,h}, a + v_{s,h} - n_{s,h})$
    update $R_{s,h} \sim \text{Be}(b + r_{s,h}, b + v_{s,h} - n_{s,h} - r_{s,h})$
  end for
end for
update a from (14);
update b from (15).

4 Simulation study

We compared our msBP method to standard Bayesian nonparametric techniques including DP location-scale mixtures of Gaussians, DP mixtures of Bernstein polynomials, and mixtures of Polya trees, all using the R package `DPpackage`. In addition, we implemented a frequentist wavelet density estimator using the package `WaveThresh`, and a simple frequentist kernel estimator. Several simulations have been run under different simulation settings leading to qualitatively similar results. We report the results for four scenarios. Scenario 1 simulated data from a mixture of betas, 0.6Be(3, 3) + 0.4Be(21, 5); Scenario 2 used a mixture of Gaussians, $0.5N(0, 4) + 0.3N(2, 1) + 0.2N(1.5, 0.25)$; Scenario 3 generated data from a density supported on the positive real line, a mixture of a gamma and a left truncated normal, $0.9\text{Ga}(2, 2) + 0.1N_{\text{LT}}(4, 0.4)$; finally, Scenario 4 generated data from a symmetric density with two spiky modes, $0.7N(0, 4) + 0.1N(0.5, 0.01) + 0.2N(1.5, 0.4)$.

For each case, we generated sample sizes of $n = 25, 50, 100$. Each of the approaches were applied to 200 replicated data sets under each scenario. The methods were compared based on a Monte Carlo approximation to the mean Kolmogorov-Smirnov distance (KS), L_1 and L_2 distances.

To implement Algorithm 3, we exploit the binary tree structure of our modelling framework using efficient C++ code embedded into R functions. In implementing the Gibbs sampler, the first 1,000 iterations were discarded as a burn-in and the next 2,000 samples were used to calculate the posterior mean of the density on a fine grid of points. To center our prior, using a default empirical Bayes approach, we set g_0 equal to a kernel estimate. For the hyperparameters we fixed $b = 1$ and let $a \sim \text{Ga}(5, 0.5)$. We truncated the depth of the binary tree to the sixth scale. The values of the density for a wide variety of points in the domain were monitored to gauge rates of apparent convergence and mixing. The trace plots showed excellent mixing, and the Geweke (1992) diagnostic suggested rapid convergence.

Table 1: Mean Kolmogorov Smirnoff (KS) distance, mean L_1 distance (L_1), and mean L_2 distance (L_2) between the true densities and the posterior msBP estimate (msBP), posterior DP mixture of Gaussians estimate (DPM), posterior DP mixture of Bernstein Polynomials estimate (DPB), posterior Polya's Tree estimate (PT), frequentist wavelet estimate (W), and frequantist kernel smoothing estimate (K) for Scenario 1 (S1), Scenario 2 (S2), Scenario 3 (S3), and Scenario 4 (S4). Mean distances computed over 200 samples, with Monte Carlo error in parenthesis

		$n = 25$			$n = 50$			$n = 100$		
		KS	L_1	L_2	KS	L_1	L_2	KS	L_1	L_2
S1	msBP	0.9616 (0.28)	15.3337 (3.79)	9.0286 (4.06)	0.8529 (0.20)	12.4909 (2.78)	5.8835 (2.59)	0.7318 (0.20)	10.2247 (2.44)	4.1602 (1.96)
	DPM	1.5785 (0.16)	18.1684 (1.78)	15.659 (2.76)	1.4137 (0.15)	18.1139 (1.50)	13.4228 (2.39)	1.3558 (0.17)	18.2278 (1.47)	13.0673 (2.46)
	DPBP	1.2443 (0.19)	22.6341 (2.42)	15.9829 (3.53)	0.9245 (0.27)	15.3053 (3.83)	8.2186 (4.03)	0.6147 (0.24)	9.7378 (3.12)	3.3916 (2.16)
	PT	2.4917 (0.00)	952.1645 (2.14)	1391.3997 (2.74)	2.4917 (0.01)	951.1084 (1.26)	1389.2295 (1.43)	2.4917 (0.01)	951.5270 (0.94)	1389.7410 (1.11)
	W	1.6867 (0.05)	26.5373 (0.81)	23.2277 (1.27)	1.6481 (0.04)	25.8640 (0.73)	22.1622 (1.15)	1.6425 (0.03)	25.7625 (0.54)	21.9891 (0.84)
	K	1.0629 (0.24)	15.8933 (3.44)	9.4448 (3.47)	0.8812 (0.21)	12.6056 (2.78)	6.0769 (2.70)	0.7623 (0.19)	10.3960 (2.53)	4.3419 (1.99)
S2	msBP	0.0947 (0.03)	1.5028 (0.32)	0.0812 (0.03)	0.0742 (0.02)	1.1060 (0.26)	0.0441 (0.02)	0.0642 (0.01)	0.9616 (0.18)	0.0327 (0.01)
	DPM	0.1385 (0.06)	1.7884 (0.53)	0.1389 (0.08)	0.1012 (0.04)	1.3192 (0.40)	0.0728 (0.04)	0.0700 (0.03)	0.9485 (0.30)	0.0372 (0.02)
	DPBP	0.2339 (0.01)	4.3513 (0.05)	0.6461 (0.01)	0.2339 (0.01)	4.4880 (0.07)	0.6648 (0.01)	0.2339 (0.01)	4.5672 (0.07)	0.6783 (0.01)
	PT	0.2347 (0.01)	94.2408 (0.44)	13.9915 (0.06)	0.2339 (0.01)	93.9891 (0.33)	13.9568 (0.03)	0.2339 (0.01)	93.8067 (0.28)	13.9393 (0.02)
	W	0.1424 (0.05)	2.1501 (0.66)	0.1756 (0.10)	0.1027 (0.03)	1.5620 (0.44)	0.0917 (0.05)	0.0717 (0.02)	1.1410 (0.31)	0.0468 (0.02)
	K	0.0931 (0.02)	1.4714 (0.31)	0.0767 (0.03)	0.0778 (0.02)	1.1730 (0.26)	0.0485 (0.02)	0.0665 (0.02)	0.9893 (0.18)	0.0344 (0.01)
S3	msBP	0.2806 (0.05)	2.7758 (0.77)	0.3854 (0.18)	0.2571 (0.04)	2.2984 (0.64)	0.2770 (0.13)	0.2252 (0.03)	1.8722 (0.43)	0.1907 (0.07)
	DPM	0.2494 (0.07)	2.8651 (0.70)	0.3922 (0.20)	0.2276 (0.06)	2.3452 (0.58)	0.2760 (0.14)	0.1938 (0.05)	1.8194 (0.31)	0.1735 (0.07)
	DPBP	0.5137 (0.04)	6.8264 (0.21)	2.1555 (0.16)	0.5735 (0.03)	7.0762 (0.21)	2.4045 (0.15)	0.6019 (0.01)	7.1933 (0.20)	2.5392 (0.13)
	PT	0.6621 (0.01)	157.7443 (0.85)	65.6996 (0.31)	0.6621 (0.01)	157.2414 (0.50)	65.5554 (0.18)	0.6621 (0.01)	156.9909 (0.26)	65.4821 (0.10)
	W	0.2982 (0.05)	3.4876 (0.71)	0.4979 (0.17)	0.2759 (0.04)	3.1490 (0.41)	0.4145 (0.10)	0.2599 (0.02)	2.9623 (0.25)	0.3631 (0.05)
	K	0.2802 (0.05)	2.963 (0.92)	0.4318 (0.23)	0.2521 (0.04)	2.4200 (0.69)	0.3006 (0.14)	0.2231 (0.03)	1.9428 (0.45)	0.2042 (0.08)
S4	msBP	0.2942 (0.04)	4.3193 (0.82)	0.5608 (0.12)	0.2943 (0.03)	3.6779 (0.57)	0.5092 (0.05)	0.2856 (0.02)	3.4838 (0.35)	0.4759 (0.04)
	DPM	0.3203 (0.06)	5.0048 (0.80)	0.7094 (0.25)	0.3037 (0.05)	4.4272 (0.63)	0.5958 (0.17)	0.2966 (0.04)	3.9836 (0.59)	0.5428 (0.15)
	DPBP	0.4995 (0.01)	8.9148 (0.16)	1.8019 (0.05)	0.4995 (0.01)	9.0004 (0.13)	1.7803 (0.05)	0.4995 (0.01)	9.0851 (0.08)	1.7881 (0.03)
	PT	0.4995 (0.01)	93.1303 (0.78)	20.5193 (0.17)	0.4995 (0.01)	92.7538 (0.65)	20.4479 (0.17)	0.4995 (0.01)	92.4458 (0.52)	20.4152 (0.14)
	W	0.2990 (0.06)	5.4053 (0.76)	0.7075 (0.23)	0.2831 (0.04)	4.613 (0.53)	0.5752 (0.12)	0.2734 (0.03)	4.0647 (0.42)	0.5036 (0.08)
	K	0.3000 (0.04)	4.3220 (0.83)	0.5834 (0.14)	0.2924 (0.03)	3.799 (0.61)	0.5143 (0.08)	0.2834 (0.02)	3.5222 (0.44)	0.4744 (0.05)

Note: 0.00 stands for "< 0.01"

The results of the simulation are reported in Table 1 and Figure 3. The proposed method performs better or equally to the best competitor in almost all scenarios and sample sizes. The worst performance in each case is obtained for mixtures of Polya trees, with overly-spiky density estimates leading to higher distances from the truth. In Scenario 1 the msBP approach beats all the competitors, except in large sample sizes when single-scale DP mixtures of Bernstein polynomials are comparable. In Scenario 2 the msBP approach is comparable to the frequentist kernel smoother estimator. In scenario 3 the msBP approach is comparable to DP location-scale mixtures and finally, in Scenario 4 our multiscale approach is clearly performing better than any other method.

5 Extensions

An appealing aspect of the proposed method is ease of generalization to include predictors, hierarchical dependence, time series, spatial structure and so on. To incorporate additional structure, one can replace model (2) for the stopping and right path probabilities with an appropriate variant. Similar extensions have been proposed for single resolution mixture models by replacing the beta random variables in a stick-breaking construction with probit regressions (Chung and Dunson 2009; Rodriguez and Dunson 2011), logistic regressions (Ren et al. 2011) or broader stochastic processes (Pati et al. 2013). We focus here on one interesting extension to the under-studied problem of Bayesian multiscale inferences on differences between groups.

5.1 Multiscale testing of group differences

Motivated by epigenetic data, we propose Bayesian multiscale hypothesis tests of group differences using multiscale Bernstein polynomials. DNA methylation arrays collect data on epigenetic modifications at a large number of CpG sites. Let $y_i = (y_{i1}, \ldots, y_{ip})^T$ denote the DNA methylation data for patient i at p different sites, with $d_i \in \{0, 1\}$ denoting the patient's disease status, either $d_i = 0$ for controls or $d_i = 1$ for cases. Current standard analyses rely on independent screening using t-tests to assess differences between cases and control at each site. However, DNA methylation data are constrained to $y_{ij} \in (0, 1)$ and tend to have a complex distribution having local spikes and varying smoothness.

As illustration we focus on nonparametric independent screening; the approach is easily adapted to accommodate dependence across sites. We center our prior on the uniform as a default. The density of y_{ij} given $d_i = 0$ is modeled as in previous sections. Let $H_0 : f_0 = f_1$ denote the *global* null hypothesis of no difference between groups, with $H_1 : f_0 \neq f_1$ denoting the alternative. Using an msBP representation, $f_0 = f_1$ if the groups share weights over the dictionary of beta densities. If $f_0 \neq f_1$, we may have the same weights on the dictionary elements up to a given scale, so that the densities are equivalent up to that scale but not at finer scales. With this in mind, let $H_0^s : f_0^s = f_1^s$ denote the null hypothesis of no differences between groups at scale s, and $H_1^s : f_0^s \neq f_1^s$ the alternative. As H_0^0 is true with probability one, we set $S_{0,1} = 0$ and concentrate on H_0^s for $s \geq 1$.

Each of the n subjects in the sample takes a path through the binary tree, stopping at a finite depth. Let $\mathcal{I}^s = \{i : s_i \geq s\}$ index the subjects *surviving* up to scale s and

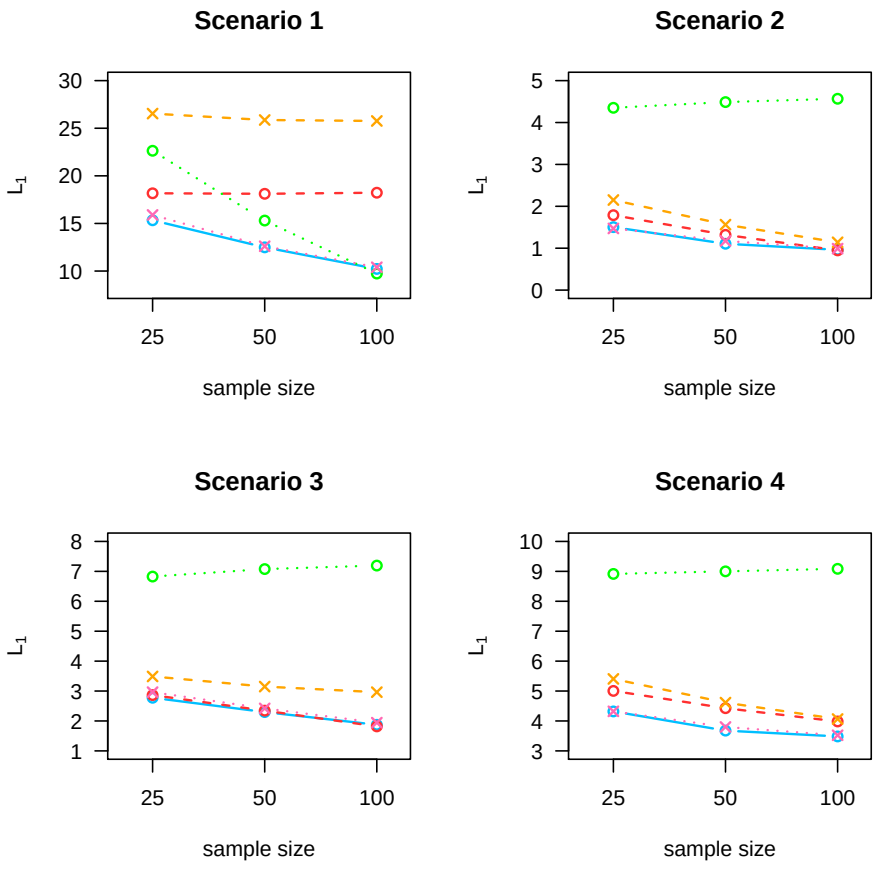

Figure 3: Mean L_1 distance between the true densities and the posterior msBP estimate (continuous line, circle dots), posterior DP mixture of Gaussians estimate (dashed line, circle dots), posterior DP mixture of Bernstein polynomials estimate (dotted line, circle dots), frequentist wavelet estimate (dashed line, "×" dots), and frequentist kernel smoothing estimate (dotted line, "×" dots) under the four scenarios. The posterior mixture of Polya trees estimate is far away and it is not reported for graphical reasons.

let $\mathcal{N}^s$ denote the actions of these subjects at scale s, including stopping or progressing downward to the left or right for each of the nodes. Subscripts (d) on $\mathcal{I}^s$ and $\mathcal{N}^s$ denote the restriction to subjects having $d_i = d$. Conditionally on H_0^s, the probabilities for each scale s action are the same in the two groups and the likelihood of actions $\mathcal{N}^s$ is

$$\begin{aligned}
\text{pr}(\mathcal{N}^s|H_0^s) &= \int_{\mathcal{T}} \text{pr}(\mathcal{N}^s|\mathcal{T})\text{pr}(\mathcal{T}|a,b)d\mathcal{T} \\
&= \left\{\frac{\Gamma(a+1)}{\Gamma(a)}\frac{\Gamma(2b)}{\Gamma(b)^2}\right\}^{2^s} \int_{\mathcal{T}} \prod_{h=1}^{2^s} S_{s,h}^{n_{s,h}}(1-S_{s,h})^{\hat{a}_{s,h}-1}R_{s,h}^{\hat{b}_{s,h}-1}(1-R_{s,h})^{\hat{c}_{s,h}-1}d\mathcal{T} \\
&= \left\{\frac{\Gamma(a+1)\Gamma(2b)}{\Gamma(a)\Gamma(b)^2}\right\}^{2^s} \prod_{h=1}^{2^s} \frac{\Gamma(1+n_{s,h})\Gamma(\hat{a})}{\Gamma(a+v_{s,h}+1)}\frac{\Gamma(\hat{b})\Gamma(\hat{c})}{\Gamma(2b+v_{s,h}-n_{s,h})}, \qquad (16)
\end{aligned}$$

where $\hat{a}_{s,h} = a + v_{s,h} - n_{s,h}$, $\hat{b}_{s,h} = b + r_{s,h}$, and $\hat{c}_{s,h} = b + v_{s,h} - n_{s,h} - r_{s,h}$. Similarly under H_1 we have

$$\begin{aligned}
\text{pr}(\mathcal{N}^s|H_1^s) =& \text{pr}(\mathcal{N}_{(0)}^s|H_1^s) \times \text{pr}(\mathcal{N}_{(1)}^s|H_1^s) \\
&= \left\{\frac{\Gamma(a+1)\Gamma(2b)}{\Gamma(a)\Gamma(b)^2}\right\}^{2^{2s}} \prod_{h=1}^{2^s} \frac{\Gamma(1+n_{s,h}^{(0)})\Gamma(\hat{a}^{(0)})}{\Gamma(a+v_{s,h}^{(0)}+1)}\frac{\Gamma(\hat{b}^{(0)})\Gamma(\hat{c}^{(0)})}{\Gamma(2b+v_{s,h}^{(0)}-n_{s,h}^{(0)})} \times \\
&\prod_{h=1}^{2^s} \frac{\Gamma(1+n_{s,h}^{(1)})\Gamma(\hat{a}^{(1)})}{\Gamma(a+v_{s,h}^{(1)}+1)}\frac{\Gamma(\hat{b}^{(1)})\Gamma(\hat{c}^{(1)})}{\Gamma(2b+v_{s,h}^{(1)}-n_{s,h}^{(1)})}, \qquad (17)
\end{aligned}$$

where $v_{s,h}^{(d)}$ is the number of subjects passing through node (s,h) in group d, $n_{s,h}^{(d)}$ is the number of subjects stopping at node (s,h) in group d, and $r_{s,h}^{(d)}$ is the number of subjects that continue to the right after passing through node (s,h) in group d, with $d = 0, 1$.

Combining (16)–(17) we can obtain a closed form for the posterior probability of H_0 being true at scale s, given $\mathcal{N}_{(0)}^s$ and $\mathcal{N}_{(1)}^s$:

$$\text{pr}(H_0^s|\mathcal{N}_{(0)}^s, \mathcal{N}_{(1)}^s) = \frac{P_0^s \text{pr}(\mathcal{N}_{(0)}^s, \mathcal{N}_{(1)}^s|H_0^s)}{P_0^s \text{pr}(\mathcal{N}_{(0)}^s, \mathcal{N}_{(1)}^s|H_0^s) + (1-P_0^s)\text{pr}(\mathcal{N}_{(0)}^s, \mathcal{N}_{(1)}^s|H_1^s)}, \qquad (18)$$

where P_0^s is our prior guess for the null being true at scale s. The global null will be the cumulative product of the $\text{pr}(H_0^s|\mathcal{N}_{(0)}^s, \mathcal{N}_{(1)}^s)$ for each scale. An interesting feature of this formulation is to have a multiscale hypothesis testing setup. Indeed the posterior probability of H_0 up to scale $\tilde{s}$ will be $\prod_{s \le \tilde{s}} \text{pr}(H_0^s|\mathcal{N}_{(0)}^s, \mathcal{N}_{(1)}^s)$ and hence the hypothesis that two groups have the same distribution may have high posterior probability for coarse scales, but can be rejected for a finer scale.

5.2 Posterior computation

The conditional posterior probability for H_0^s in (18) is simple, but not directly useful due to the dependence on the unknown $\mathcal{N}^s$ allocations. To marginalize out these allocations, we modify Algorithm 3. For node h at scale s, let $\pi_{s,h}^{(0)}$ denote the weight under H_0^s and $\pi_{s,h}^{(1,d)}$ for $d = 0, 1$ denote the group-specific weights under H_1^s. At each

iteration, the allocation of subject i of group d will be made according to the tree of weights given by

$$\pi_{s,h}^{(d)} = P(H_0^s|\mathcal{N}_{(0)}^s, \mathcal{N}_{(1)}^s)\pi_{s,h}^{(0)} + \{1 - P(H_0^s|\mathcal{N}_{(0)}^s, \mathcal{N}_{(1)}^s)\}\pi_{s,h}^{(1,d)}. \quad (19)$$

Given the allocation one can calculate all the quantities in (16)–(17) and then update the stopping and descending probabilities under H_0 and H_1 following (13) and the posterior of the null following (18) up to a desired upper scale.

6 Application

We illustrate our approach on a methylation array dataset for $n = 597$ breast cancer samples registered at $p = 21{,}986$ CpG sites (Cancer Genome Atlas Network 2012). We test for differences between tumors that are identified as basal-like ($n_0 = 112$) against those that are not ($n_1 = 485$) at each CpG site. This same problem was considered in a single scale manner by Lock and Dunson (2014) using finite mixtures of truncated Gaussians.

We run the Gibbs sampler reported in Algorithm 4 in the Appendix, assuming a uniform prior for P_0^s for each scale s. We fixed the maximum scale to 4 as an upper bound, as finer scale tests were not thought to be interpretable. The sampler is run for 2,000 iterations after 1,000 burn-in iterations. The chains mix well and converge quickly for all sites and all scales.

The posterior distribution of $1-P_0^s$ for each scale provides a summary of the overall proportion of CpG sites for which there was a difference between the two groups. The estimated posterior means for these probabilities were 0.04, 0.07, 0.05 and 0.03, respectively, for scales $1, \ldots, 4$. This suggests that DNA methylation levels were different for a small minority of the CpG sites, which is as expected. Examining the posterior probabilites of H_1^s across the 21,986 CpG sites, consistently with the estimates for $1 - P_0^s$, we find that scale-specific estimated posterior probabilities are close to zero for most sites. Focusing on the 1,696 sites for which the overall posterior probability of H_1 is greater than 0.5, we calculated the minimal scale showing evidence of a difference, $\min\{s : \hat{Pr}(H_1^s|-) > 0.5\}$, with $\hat{Pr}(H_1^s|-)$ denoting the estimated posterior probability. The proportions of sites having minimal scale equal to $1, 2, 3, 4$ were $47\%, 43\%, 7\%, 3\%$ respectively.

Figure 4 shows $\hat{Pr}(H_1^s|-)$ for these 1,696 sites. In the top right quadrant we report those sites having minimal scale equal to 1. Two different patterns are evident: (1) consistently high $\hat{Pr}(H_1^s|-)$, with differences evident at the coarse scale. Site *cg00117172* is among those and its sample distribution is reported in panel (a) of Figure 5. (2) moderate $\hat{Pr}(H_1^s|-)$ for $s = 1$, with clear evidence at $s = 2$. Averages of the sites in these two groups are shown with thick dashed lines.

The top right panel, representing sites having minimal scale equal to 2, presents two patterns: (1) no differences at scale one but clear evidence of H_1 at scale two. Site *cg00186954* in panel (b) of Figure 5 has this behavior. (2) moderately growing evidence for H_1 for increasing scale level. The bottom two panels show results for sites having minimal scale equal to 3 and 4, showing again two different patterns: (1) A group with mild or no evidence for H_1 up to scale 3 and 4, respectively (e.g. site *cg20603888* reported in panel (c) of Figure 5), and (2) another group with increasing

evidence for increasing scale. These scale-specific significant tests are interesting in that coarser scale differences are more likely to be biologically significant, while very fine scale differences may represent local changes with minor impact.

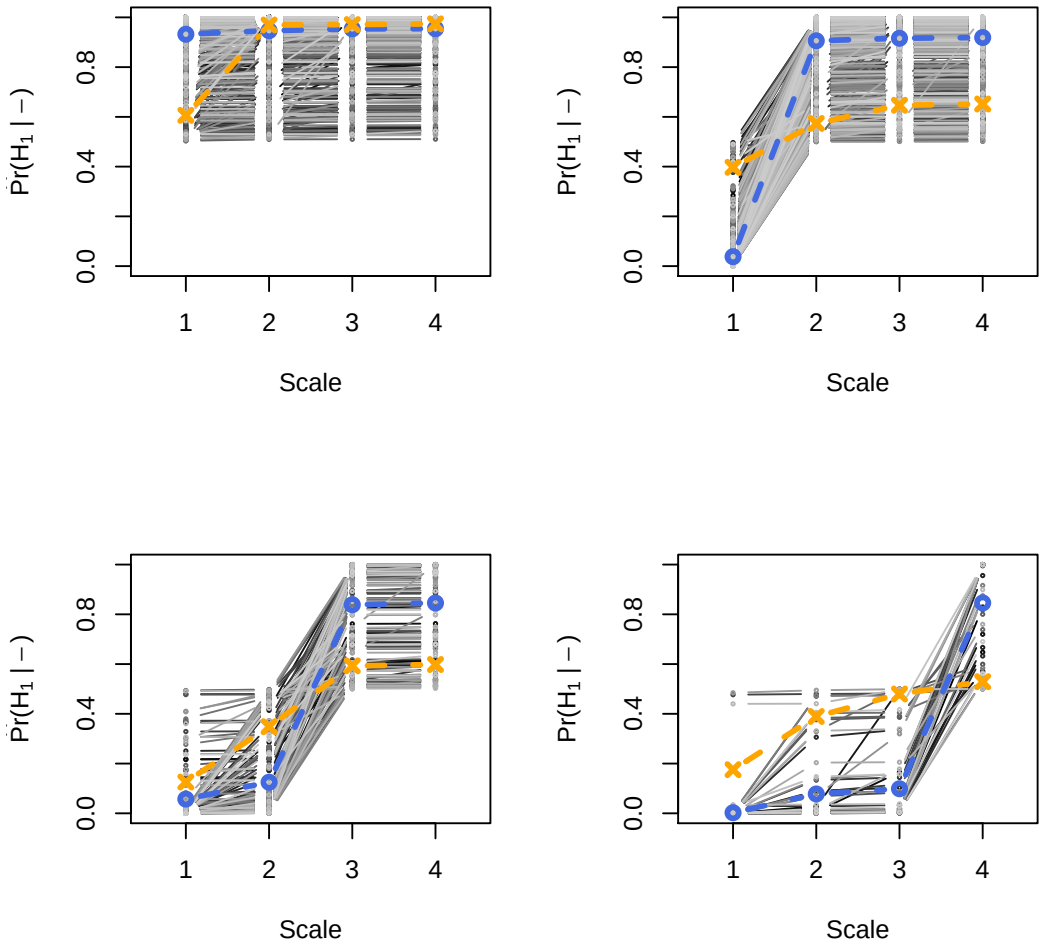

Figure 4: Posterior mean probabilities of H_1 depending on scale for the 1,696 sites, with some evidence of differences in the two groups, grouped in subplots by minimal scale showing $\hat{Pr}(H_1^s|-) > 0.5$ for $s = 1, \ldots 4$. Within each panel, the thick dashed lines represents the average between the sites in two clusters showing different patterns.

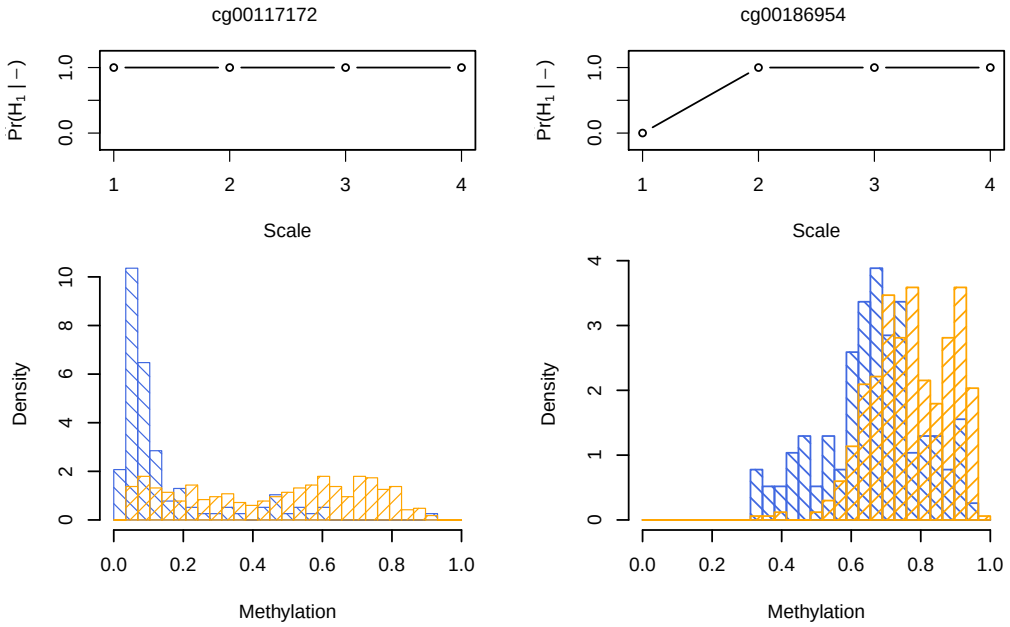

(a)

(b)

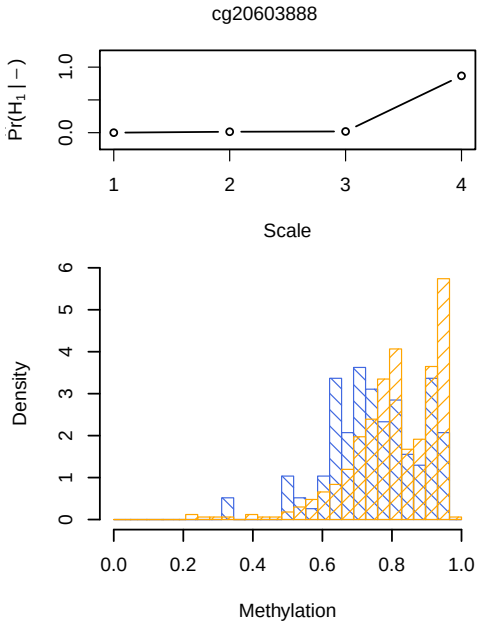

(c)

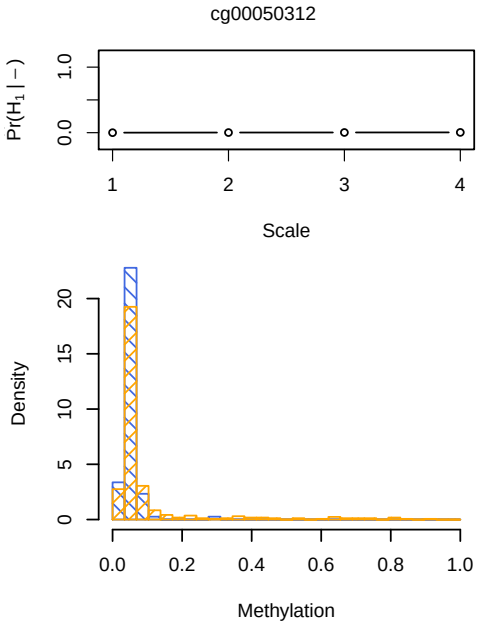

(d)

Figure 5: Histogram of the methylation for the basal (decreasing 45 degree angle shading) and non-basal (increasing 45 degree angle shading) samples for four CpG sites and posterior mean probabilities of H_1 in function of scale.

7 Discussion

Existing Bayesian nonparametric multiscale tools for density estimation have unappealing characteristics, such as favoring overly-spiky densities. Our framework overcomes such limitations. We have demonstrated some practically appealing properties, including simplicity of formulation and ease of computation, and proposed an extension for Bayesian multiscale hypothesis testing of group differences. Multiscale hypothesis testing is of considerable interest in itself, and provides a new view on the topic of nonparametric testing of group differences, with many interesting facets. For example, it can be argued that in large samples there will always be small local differences in the distributions between groups, which may not be scientifically relevant. By allowing scale-specific tests, we accommodate the possibility of focusing inference on the range of relevant scales in an application, providing additional insight into the nature of the differences. We also accommodate scale-specific adaptive borrowing of information across groups in density estimation; extensions to include covariates and hierarchical structure are straightforward.

Acknowledgement

The authors thanks Eric Lock for helpful comments on Section 5 and Roberto Vigo for comments on the code implementation.

Appendix

Proof of Lemma 1. For finite N define $\Delta_N = 1 - \sum_{s=0}^{N} \sum_{h=1}^{2^s} \pi_{s,h}$, for which the following inequality holds:

$$\Delta_N = \sum_{h=1}^{2^N} \prod_{r \leq N} (1 - S_{r,g_{Nhr}}) T_{r-1,g_{Nhr}} \leq 2^N \max_{h=1,\ldots,2^N} \prod_{r \leq N} (1 - S_{r,g_{Nhr}}) T_{r-1,g_{Nhr}}. \quad (20)$$

To establish (5), it is sufficient to take the limit of Δ_N for $N \to \infty$ and show that it converges to 0 a.s. To this end, take the logarithm of the right hand side of (20),

$$\log(\Delta_N) \leq \max_{h=1,\ldots,2^N} \sum_{r \leq N} \log \left\{ 2^N (1 - S_{r,g_{Nhr}}) T_{r-1,g_{Nhr}} \right\}, \quad (21)$$

and notice that for each $h = 1, \ldots, 2^N$ we have

$$E \left\{ 2^N (1 - S_{r,g_{Nhr}}) T_{r-1,g_{Nhr}} \right\} = 2^N \left(\frac{a}{a+1} \right) \frac{1}{2^N} = \frac{a}{a+1}. \quad (22)$$

Therefore taking $N \to \infty$, by Kolmogorov's three series theorem and Jensen's inequality, the argument of the maximum of (21), converges to $-\infty$ a.s. for each h. Thus Δ_N converges to 0 a.s. which concludes the proof. □

Detail on moments of $F(A)$. The expectation of $F(A)$ is simply

$$\begin{aligned}
E[F(A)] &= E\left[\sum_{s=0}^{\infty}\sum_{h=1}^{2^s}\pi_{s,h}\int_A \mathrm{Be}(y;h,2^s-h+1)\right] \\
&= \sum_{s=0}^{\infty}\frac{1}{1+a}\left(\frac{a}{1+a}\right)^s\frac{1}{2^s}\sum_{h=1}^{2^s}\int_A \mathrm{Be}(y;h,2^s-h+1) \\
&= \sum_{s=0}^{\infty}\frac{1}{1+a}\left(\frac{a}{1+a}\right)^s\lambda(A) \\
&= \lambda(A)\sum_{s=0}^{\infty}\frac{1}{1+a}\left(\frac{a}{1+a}\right)^s = \lambda(A),
\end{aligned}$$

where the third equality follows from the fact that the average measure over scale s beta dictionary densities of any region A equals the Lebesgue measure of A. □

Proof of Lemma 2. First note that twice the total variation distance between two measures P^s and P equals the L_1 distance between the densities f^s and f. For the expectation, the following holds

$$E\left[\int\left|f^s(y)-f(y)\right|dy\right] = \int E\left[\left|f^s(y)-f(y)\right|\right]dy$$

by Fubini's theorem. Now since

$$\left|f^s(y)-f(y)\right| = f^s(y)-f(y)+2\max\{f(y)-f^s(y),0\},$$

it is sufficient to prove that the expectation of $f^s(y)-f(y)$ is null. This can be done, noting that for each $y\in[0,1]$ and for each scale s, the quantity $2^{-s}\sum_{h=1}^{2^s}\mathrm{Be}(y;h,2^s-h+1)=1$. Hence

$$\begin{aligned}
&\sum_{h=1}^{2^s}E[\tilde{\pi}_{s,h}]\mathrm{Be}(y;h,2^s-h+1)-\sum_{l=s}^{\infty}\sum_{h=1}^{2^l}E[\pi_{l,h}]\mathrm{Be}(y;h,2^l-h+1) = \\
&\quad= \sum_{h=1}^{2^s}E[\tilde{\pi}_{s,h}-\pi_{s,h}]\mathrm{Be}(y;h,2^s-h+1)-\sum_{l=s+1}^{\infty}\sum_{h=1}^{2^l}E[\pi_{l,h}]\mathrm{Be}(y;h,2^l-h+1) \\
&\quad= \left(\frac{a}{1+a}\right)^{s+1}\frac{1}{2^s}\sum_{h=1}^{2^s}\mathrm{Be}(y;h,2^s-h+1)-\sum_{l=s+1}^{\infty}\frac{1}{1+a}\left(\frac{a}{1+a}\right)^l\frac{1}{2^l}\sum_{h=1}^{2^l}\mathrm{Be}(y;h,2^l-h+1) \\
&\quad= \left(\frac{a}{1+a}\right)^{s+1}-\sum_{l=s+1}^{\infty}\frac{1}{1+a}\left(\frac{a}{1+a}\right)^l = \left(\frac{a}{1+a}\right)^{s+1}-\left(\frac{a}{1+a}\right)^{s+1}=0,
\end{aligned}$$

which concludes the first part of proof. Now consider

$$
\begin{aligned}
&\int \left| \sum_{l=0}^{s} \sum_{h=1}^{2^l} \tilde{\pi}_{l,h} \mathrm{Be}(y; h, 2^s - h + 1) - \sum_{l=0}^{\infty} \sum_{h=1}^{2^l} \pi_{l,h} \mathrm{Be}(y; h, 2^l - h + 1) \right| dy \\
&= \int \left| \sum_{l=0}^{\infty} \sum_{h=1}^{2^l} (\tilde{\pi}_{l,h} - \pi_{l,h}) \, \mathrm{Be}(y; h, 2^s - h + 1) \right| dy \\
&\leq \int \sum_{l=0}^{\infty} \sum_{h=1}^{2^l} |(\tilde{\pi}_{l,h} - \pi_{l,h}) \, \mathrm{Be}(y; h, 2^s - h + 1)| \, dy \\
&= \sum_{l=0}^{\infty} \sum_{h=1}^{2^l} |(\tilde{\pi}_{l,h} - \pi_{l,h})| \int \mathrm{Be}(y; h, 2^s - h + 1) dy = \sum_{l=0}^{\infty} \sum_{h=1}^{2^l} |(\tilde{\pi}_{l,h} - \pi_{l,h})| \,,
\end{aligned}
$$

where the inequality holds since for each y the absolute values of the sum is less than the sum of the absolute values. Since the first moment is null the variance is

$$
\begin{aligned}
E\left[\left\{ \int |f^s(y) - f(y)| \, dy \right\}^2 \right] &= E\left[\left(\sum_{l=0}^{\infty} \sum_{h=1}^{2^l} |\tilde{\pi}_{l,h} - \pi_{l,h}| \right)^2 \right] \\
&= E\left[\left(\sum_{h=1}^{2^s} |\tilde{\pi}_{s,h} - \pi_{s,h}| + \sum_{l=s+1}^{\infty} \sum_{h=1}^{2^l} \pi_{s,h} \right)^2 \right] \\
&\leq 2E\left[\left(\sum_{h=1}^{2^s} |\tilde{\pi}_{s,h} - \pi_{s,h}| \right)^2 + \left(\sum_{l=s+1}^{\infty} \sum_{h=1}^{2^l} \pi_{s,h} \right)^2 \right].
\end{aligned}
$$

We study separately the expecations of the two summands above. For each $h = 1 \ldots, 2^s$, $\tilde{\pi}_{s,h} \geq \pi_{s,h}$, thus the fist expectation is

$$
\begin{aligned}
E\left\{ \left(\sum_{h=1}^{2^s} \tilde{\pi}_{s,h} - \pi_{s,h} \right)^2 \right\} &\leq E\left\{ \left(\sum_{h=1}^{2^s} \tilde{\pi}_{s,h} \right)^2 + \left(\sum_{h=1}^{2^s} \pi_{s,h} \right)^2 \right\} \\
&\leq E\left(\sum_{h=1}^{2^s} \tilde{\pi}_{s,h} + \sum_{h=1}^{2^s} \pi_{s,h} \right) \\
&= \left(\frac{a}{1+a} \right)^s + \frac{1}{1+a} \left(\frac{a}{1+a} \right)^s,
\end{aligned}
$$

where the first inequality holds removing twice the cross product, and the second since the quantities are strictly less than one. The second expectation is simply

$$
E\left\{ \left(\sum_{l=s+1}^{\infty} \sum_{h=1}^{2^l} \pi_{s,h} \right)^2 \right\} \leq E\left(\sum_{l=s+1}^{\infty} \sum_{h=1}^{2^l} \pi_{s,h} \right) = \left(\frac{a}{1+a} \right)^{s+1}.
$$

It follows that the variance is less than $2\{a/(1+a)\}^s$, that concludes the proof. □

Algorithm 4 Gibbs sampler steps for posterior computation for multiscale hypothesis testing of group differences using msBP prior

for $j = 1, \ldots, p$ **do**
  Compute the threes for the node allocation according to (19).
  for $i = 1, \ldots, n$ **do**
    assign observation i at site j to a cluster (s_i, h_i) as in Algorithm 2 using the tree of weights of last step
  end for
  compute $n_{s,h}$, $v_{s,h}$, and $r_{s,h}$;
  compute $n_{s,h}^{(j)}$, $v_{s,h}^{(j)}$, and $r_{s,h}^{(j)}$ for $j = 0, 1$;
  let s_{MAX} the maximum occupied scale;
  for $s = 0, \ldots, s_{\text{MAX}}$ **do**
    for $h = 1, \ldots, 2^s$ **do**
      update $S_{s,h} \sim \text{Be}(1 + n_{s,h}, a + v_{s,h} - n_{s,h})$, $R_{s,h} \sim \text{Be}(b + r_{s,h}, b + v_{s,h} - n_{s,h} - r_{s,h})$
      update $S_{s,h}^{(0)} \sim \text{Be}(1 + n_{s,h}^{(0)}, a + v_{s,h}^{(0)} - n_{s,h}^{(0)})$, $R_{s,h}^{(0)} \sim \text{Be}(b + r_{s,h}^{(0)}, b + v_{s,h}^{(0)} - n_{s,h}^{(0)} - r_{s,h}^{(0)})$
      update $S_{s,h}^{(1)} \sim \text{Be}(1 + n_{s,h}^{(1)}, a + v_{s,h}^{(1)} - n_{s,h}^{(1)})$, $R_{s,h}^{(1)} \sim \text{Be}(b + r_{s,h}^{(1)}, b + v_{s,h}^{(1)} - n_{s,h}^{(1)} - r_{s,h}^{(1)})$
    end for
  end for
  compute the trees of weights under H_0 and H_1 for the two groups
  for $s = 0, \ldots, s_{\text{MAX}}$ **do**
    compute $P_m^s = \text{pr}(H_0^s | \mathcal{N}_{(0)}^s, \mathcal{N}_{(1)}^s)$ as in (18).
  end for
end for
Draw $P_0^s \sim \text{Be}(1 + \sum_{m=1}^M P_m^s, 1 + M - \sum_{m=1}^M P_m^s)$

References

Abramovich, F., Sapatinas, T., and Silverman, B. W. (1998). Wavelet thresholding via a bayesian approach. *Journal of the Royal Statistical Society, Series B: Statistical Methodology*, 60:725–749.

Adams, R. P., Z., G., and I., J. M. (2010). Tree-structured stick breaking for hierarchical data. *Advances in Neural Information Processing Systems*, 23:19–27.

Cancer Genome Atlas Network (2012). Comprehensive molecular portraits of human breast tumours. *Nature*, 460:61–70.

Chen, G., Iwen, M., Chin, S., and Maggioni, M. (2012). A fast multiscale framework for data in high-dimensions: Measure estimation, anomaly detection and compressive measurements. *IEEE Visual Communications and Image Processing*, pages 1–12.

Chung, Y. and Dunson, D. (2009). Nonparametric bayes conditional distribution modeling with variable selection. *Journal of the American Statistical Association*, 104:1646–1660.

Clyde, M. and George, E. I. (2000). Flexible empirical bayes estimation for wavelets. *Journal of the Royal Statistical Society, Series B: Statistical Methodology*, 62:681–698.

Clyde, M., Parmigiani, G., and Vidakovic, B. (1998). Multiple shrinkage and subset selection in wavelets. *Biometrika*, 85:391–401.

Donoho, D., Johnstone, I., Kerkyacharian, G., and Picard, D. (1996). Density estimation by wavelet thresholding. *Annals of Statistics*, 24:508–539.

Garcia-Trevino, E. and Barria, J. (2012). Online wavelet-based density estimation for non-stationary streaming data. *Computational Statistics & Data Analysis*, 56:327–344.

Geweke, J. (1992). Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments. In Bernardo, J. M., Berger, J. O., Dawid, A. P., and Smith, A. F. M., editors, *Bayesian Statistics 4*. Oxford: Oxford University Press.

Ghosal, S. (2001). Convergence rates for density estimation with Bernstein polynomials. *The Annals of Statistics*, 29:1264–1280.

Hanson, T. and Johnson, W. O. (2002). Modeling regression error with a mixture of polya trees. *Journal of the American Statistical Association*, 97:1020–1033.

Ishwaran, H. and James, Lancelot, F. (2001). Gibbs sampling methods for stick breaking priors. *Journal of the American Statistical Association*, 96(453):161–173.

Kalli, M., Griffin, J., and Walker, S. (2011). Slice sampling mixture models. *Statistics and Computing*, 21(1):93–105.

Lavine, M. (1992a). some aspects of Polya tree distributions for statistical modelling. *Annals of Statistics*, 20:1222–1235.

Lavine, M. (1992b). more aspects of Polya tree distributions for statistical modelling. *Annals of Statistics*, 22:1161–1176.

Lock, E. F. and Dunson, D. B. (2014). Shared kernel bayesian screening. *arXive*, 1311.0307:1–20.

Locke, J. and Peter, A. (2013). Multiwavelet density estimation. *Applied Mathematics and Computation*, 219:6002–6015.

Mauldin, D., Sudderth, W. D., and Williams, S. C. (1992). Polya trees and random distributions. *Annals of Statistics*, 20:1203–1203.

Niu, S. L. (2012). Nonlinear wavelet density estimation with censored dependent data. *Mathematical Methods in the Applied Sciences*, 35:293–306.

Pati, D., Dunson, D., and Tokdar, S. (2013). Posterior consistency in conditional distribution estimation. *Journal of Multivariate Analysis*, 116:456–472.

Pensky, M. and Vidakovic, B. (1999). Adaptive wavelet estimator for nonparametric density deconvolution. *Annals of Statistics*, 27:2033–2053.

Petrone, S. (1999a). Bayesian density estimation using Bernstein polynomials. *Canadian Journal of Statistics*, 27:105–126.

Petrone, S. (1999b). Random Bernstein polynomials. *Scandinavian Journal of Statistics*, 26:373–393.

Petrone, S. and Wasserman, L. (2002). Consistency of Bernstein polynomial posteriors. *Journal of the Royal Statistical Society, Series B: Statistical Methodology*, 64:79–100.

Ren, L., Du, L., Carin, L., and Dunson, D. B. (2011). Logistic stick-breaking process. *Journal of Machine Learning Research*, 12:203–239.

Rodriguez, A. and Dunson, D. (2011). Nonparametric bayesian models through probit stick-breaking processes. *Bayesian Analysis*, 6:145–177.

Sethuraman, J. (1994). A constructive definition of Dirichlet priors. *Statistica Sinica*, 4:639–650.

Wang, X., Ray, S., and Mallick, B. K. (2007). Bayesian curve classification using wavelets. *Journal of the American Statistical Association*, 102(479):962–973.